RTP Media Congestion Avoidance Techniques (rmcat)
Room: Atlantico B
Wednesday, April 6, 10:00-12:30

WG Status & Agenda Bashing (Chairs)

Mirja: I see people nodding for an interim before Berlin
Zahed: Not sure how to summarize all the results, but it would be good to do that. Ingemar is planning in Berlin to show these. An interim, say Sunday in Berlin, would be good.

Xiaoqing: Table: evaluation test cases (wired (basic), wifi and cellular) vs candidates. For NADA, no update since Yokohama.
Mirja: More discussion time in interim?
Xiaoqing: Have not seen a need to update, too early to tell. Missing cellular test cases.
Mirja: Try to decide a couple weeks before Berlin.

draft-ietf-rmcat-eval-criteria-05 (Varun Singh, remote)

Varun: Stefan will update jitter model for 06 version of document
Varun: Any opinions moving metrics to eval-criteria (none) OK will move
Mirja: Question on Jabber (from Sergio Mena): how long are the audio and video sequences?
Varun: 500, 600 frames.
Xiaoqing: We need longer sequences.
Mo Zanaty: You can use the NetVC sequences if you want. We host them. There are a few longer ones. Netflix is adding much longer ones. Will send to the list.
Mirja: Where are these stored?
Mo: Mozilla private servers, we can mirror them to wherever.

Mirja: please post question on UDP BG traffic to mailing list. Do we need it? How do we generate it?

Mirja: Energy to work on Adaptive FEC? (Nope.)
Varun: Will let it expire, revisit later.

draft-ietf-rmcat-eval-test-03 (Zahed)

Mirja (from the floor): Cubic more appropriate than 5681, there's a draft. Take it as the default since it's a little more aggressive.
Varun: Did we already discuss this? I thought we picked this. But there was no reference at the point.
Zahed: Okay, will change to cubic.

Zahed: will move definitions of metrics to eval-criteria

Mirja: Integration with video traffic model?
Zahed: Eval-criteria gather all the information needed about the test environment, eval-test specifies a particular test case that you want to do, media source model defines synthetic media model and how to realize requirements from eval-test (section 4). Can move some more text from eval-test and point to video traffic model draft.

Xiaoqing: Here (eval-test) we have media source behavior (what). Video traffic model drafts have two references for how to generate that traffic (how).

Randell Jesup (remote): Suggest we want a test sequence with slide show content. Traffic pattern very different than live video. Null change frames if any, big spikes at slide change, typically.
Zahed: It would be good to see results from trace files, just discussion doesn't help.

Mirja: Relationship between documents: is this clear, or does there need to be alignment changes?
Xiaoqing: Updated eval-test already points to video traffic model. Need a reverse reference.
Mirja: Are you using the same terminology?
Xiaoqing: Need to realign.

draft-ietf-rmcat-wireless-tests-01 (Xiaoqing Zhu)

At slide 6

Zahed: Lots of rate jumps. Are you saturating the link? Should we change the test cases so we can see convergence?
Xiaoqing: I thought we were interested in seeing behavior at saturation. Not clear any CC can handle this well. Would be interesting for others to try.
Mirja: Ingemar just proposed to see if TCP would be any more stable in this situation. As a test to see if this test case is useful.
Xiaoqing: I don't agree that the test case is not useful if no algorithms work in the case.
Zahed: Would like to revisit if this test case should be mandatory to test, as it is now easily available.
Xiaoqing: happy to share code
Mirja: Conclusion on ML, test cases to verify algorithm is safe to run on the Internet. This may be useful as a safety case.

At slide 11

Mirja: Clarification from Varun, loss on UDP here?
Xiaoqing: Not shown, suspect yes.

At end

Mirja: From jabber, how were background flows modeled?
Xiaoqing: 5 UDP CBR flows

draft-ietf-rmcat-video-traffic-model-00 (Xiaoqing Zhu)

Mo: We (netvc) have lots of sequences available, classified by type.

On slide 7

Zahed: We have codec experts here. I haven't seen this (big frames). Is this common in production?
Mo: Scream has a model of the video encoder, this model changes the resolution at the big jumps in available bandwidth. Each model, each strategy leads to different profile. Changing frame rate is not jumpy. Changing resolution leads to large pframe or iframe, so you get big frames like this.
Mirja: from jabber, same resolution used here
Michael Ramalho: this is not typical without resolution changes

Miguel París: Do you have a test for very low frame rates? It's another different situation
Xiaoqing: That's like slide sharing, we don't have those yet.

Roni Even: Not all codecs behave the same. Some may send full frames periodically anyhow. This is all codec internal, can't predict how it will happen with different implementations

Xiaoqing: Frame rate change, do people think it is realistic to do this instead of or in addition to quantization behavior?
Roni Even: Definitely, slides will reduce frame rate.
Xiaoqing: But for conferencing content?

Randall: Default assumption to make, we should leave the frame rates constant, and adapt quantization. Preference to keep framerates high and let quantization vary, adjust resolution at the source when too far for the target. Frame rate drop is a safety valve.

Zahed: Let's not overengineer our synthetic source model. Fix on something simpler, concrete. Implementers can test variety of content in production

Mirja: Agreement on jabber to only change quantization.

At slide 9

Mo: Are these codecs operating with channel awareness?
Xiaoqing: This is a full conferencing system, hijacked the CC to fix the rate.
Mo: That brings a whole dimension of complexity. Loss feedback to a codec manifests in very different ways. Want to pick one? Pick one from a WebRTC browser.
Xiaoqing: We don't want to model codec behavior. We don't want to care about reaction in the codec.

At slide 11

Zahed: Here, resolution can change?
Xiaoqing: No
Zahed: What was the other one then?

Roni Even: Was the switch from 400->1m at the same fram

Randall: h264 was not designed as an interactive codec. You're measuring the implementation, not the algorithm, not the base codec.

Xiaoqing: People interested in gathering traces on own codec?

Zahed: Just generate with a WebRTC implementation.
Xiaoqing: People interested in that?
Zahed: Sure.

At end:

Mirja: Randall (in jabber) will try to grab traces for vp8, vp9, openh264

Michael Ramallo: On these rate changes, are these working as a token bucket? When you reset the token bucket on target change, this might account for the burst...

draft-ietf-rmcat-nada-02 (Xiaoqing Zhu)

On slide 26

Mirja: Form jabber, Why was jitter different in ns2/ns3?
Xiaoqing: Different test conditions; change in physical link rate vs UDP cross

At end

Ingemar: You might get jitter in ns3 because of interact between NADA and UDP packets in the queue?
Xiaoqing: Yes.

Zahed: More disturbances on the feedback channel? You might get cases where you cannot get 100 ms feedback delay.
Xiaoqing: We do have results...
Mirja: from Piers in Jabber, how does NADA depend on synchronized clocks?
Xiaoqing: Algorithm assumes that for queueing delay.
Mirja: And how bad is it if not sync?
Xiaoqing: Have not looked.

draft-ietf-rmcat-coupled-cc-00 (Safiqul Islam)

Mirja: This draft is ready. WGLC?
Zahed: You talked about incorporating summary results for Scream, what happened? Should it not wait for the final decision of the WG?
Mirja: No, because the way it works in the draft is these are merely examples. No problem applying general technique to Scream. Unless you want to send text, then we can wait.
Zahed: Need to think, but good to know it's ready.

draft-ietf-rmcat-sbd-04 (David Hayes)

Mirja: We'll go for experimental. Will run WGLC very soon.

Discussion on Common Feedback Format

Zahed: Implementation has its own feedback
Mirja: But Ingemar shows it's possible to map to existing formats.

Dan Romascanu: One slide had comment about overhead. Need to use either 3611 or any of the RFCs that came after, avoid this duplication of fields.

David Hayes: One statistic missing. We fall back to loss when other statistics too high.

Mirja: Candidates don't need the same information. If you look closer, it's basically the same information. Loss, bandwidth estimation, delay. Can we converge from here, does this make sense?

Xiaoqing: Quick observation: looks like feedbacks recommended by other schemes seem to be more detailed. Field values recommended in NADA can be derived from the other ones.

Stefan: (GCC) Packet identifier is mostly there for computing aggregate BW, not loss rate. Packet send time is already known at send side.

Varun: The next slide is interesting. Feedback interval is the first thing we need to converge on. Then we can decide how we can encode information, because that might depend on the interval.

Zahed: Candidate information could change based on interval?

Varun: not information, but granularity, and granularity changes with interval. Higher granularity can happen every Nth report

Mirja (from floor): We don't need to agree on one *rate*, you can request a feedback rate, but a minimum feedback rate + maximum bits for feedback leads us to encoding

Mo: Curious if those intervals are absolute, or in terms of RTT? Why would the solution candidates care about these timeframes? Media encoding times, or RTTs?

Xiaoqing: For NADA, reaction time is RTT independent. Numbers we picked is overhead/response time tradeoff.

Mo: Multiple factors, RTT, encoder reaction time, third is

Xiaoqing: Input to RTT, assume RTT below a certain level.

Mo: If someone is on a home net you don’t want to bottleneck, should be more cautions wrt feedback interval.

Xiaoqing: Source constrained.

Varun: rtcp can also get lost, complicates things.

Stefan: Feedback rate for gcc is not rtt related. Lower fb rate increases queueing delay. Connected to how the controller behaves.

Varun: I see it as a combination of RTT, media rate, and how quickly the sender can react.

Brian: Good to know what the assumed RTTs are, in any case, and how things behave outside the RTT window

Ingemar: RTT not a key design factor. Media rate more important. At low media rates not important at all.

David Hayes: SBD's time on slide 21 is dependent on packet measurements on that interval. You need enough measurements to calculate distance.
Mirja: No stability problem?
David: no

Xiaoqing: Test case draft specifies propagation delay, in NADA details as well.

Mo: Feedback means all types of feedback or rate updates? In a very low latency network, you want to repair very quickly. RMCAT should not have 100ms dropouts for media repair on local links.

Xiaoqing: Discussion is for rate adaptation only. Is loss and repair in scope? Okay to specify regular rate adaptation feedback interval, but loss intervals happen separately as soon as detected.

Mirja: How often do you anyway get RTP feedback?

Mirja: Is there an opportunity to create a common feedback format?

Zahed: I see there is convergence, I see there is need. That also means everyone has to have sender-based systems, not sure we can do that. If so, yes.

Varun: Interval still the most important bit to agree on. Good summary of what needs to go there. If we pick a feedback interval, run their CC with that interval to see if it makes sense.

Mirja: Two options: Pick feedback interval for everybody, or feedback interval negotiation.

Zahed: Negotiation might work.

Randall: In addition to negotiation... we need to agree on CC-specific negotiation, or a fixed value, or on a simple algorithm based on bandwidth, etc.

Xiaoqing: Are there any limiting factors for feedback intervals, other than overhead? Why not just pick minimum feasible?

Zahed: Negotiation helps when things are not as expected. Frequency/capacity of RTP feedback will be calculated from session bandwidth. You don't want to dedicate all of that to CC feedback.

Mo: The core of this, especially for all sender-side CCs, then you want ack and time. Keep RTPCORE and XRBLOCK in the loop so we don't reinvent any wheels.

Zahed: Here, we only discuss commonalities. New RTCP messages happen in other WGs.

Mirja: Do we already agree everything is sender-based? On Tuesday there was largely agreement, so probably yess.

Randell: Another thing to remember about high feedback rates, you're competing with media traffic.

Mirja: How do we move forward from here? Stefan's doc doesn't fit the needs yet, do we want to move forward from there?

Zahed: Makes sense, let's start from there and this slide (20).

Mirja: I propose a small design team, syncs up via telephone, bring results to mailing list. Pretty clear who has to be on the design team. Zahed will take the lead.